



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) Publication number: **0 668 565 A1**

(12)

## EUROPEAN PATENT APPLICATION

(21) Application number: **95300714.3**

(51) Int. Cl.<sup>6</sup>: **G06F 12/10**

(22) Date of filing: **06.02.95**

(30) Priority: **22.02.94 US 199098**

(43) Date of publication of application:  
**23.08.95 Bulletin 95/34**

(84) Designated Contracting States:  
**AT BE DE DK ES FR GB GR IE IT LU NL PT SE**


(71) Applicant: **ADVANCED MICRO DEVICES INC.**  
One AMD Place,  
P.O. Box 3453  
Sunnyvale, California 94088-3453 (US)

(72) Inventor: **Kranich, Uwe**  
**Arberweg 8**  
**D-85551 Heimstetten, Munich (DE)**

(74) Representative: **Wright, Hugh Ronald et al**  
**Brookes & Martin**  
**52/54 High Holborn**  
**London WC1V 6SE (GB)**

### (54) Virtual memory system.

(57) An apparatus and method for reducing the cache miss penalty in a virtual memory system is provided. The virtual memory system includes a processor core which generates virtual addresses and a cache configured to supply information in response to receipt of physical addresses. The apparatus includes a logical-to-physical translation unit which converts the virtual addresses generated by the processor core to physical addresses. The logical-to-physical translation unit includes an accurate translation unit, a speculative translation unit, and a comparing unit. The accurate translation unit accurately converts logical addresses to physical addresses. The speculative translation unit generates and transmits a speculative physical address to the cache before the accurate translation unit completes generation of the accurate physical address. When the accurate translation unit completes generation of the accurate physical address, the comparing unit compares the accurate physical address with the speculative physical address. If the accurate physical address does not match the speculative physical address, the transmission of the speculative physical address to the cache is aborted, and the accurate physical address is transmitted to the cache.



EP 0 668 565 A1

The present invention relates to a method and apparatus for decreasing the average time required to retrieve information required by a processor, and more particularly to an apparatus and method for reducing the cache miss penalty in a caching system by generating a speculative physical address in parallel with the translation of a virtual address to a physical address.

One way to increase the performance of a computer system is to decrease the time required to supply a microprocessor with the information it requires. Such information, which may include both data and instructions, is typically stored in a memory connected to the microprocessor. To access required information, the microprocessor transmits to the memory an address which corresponds to the physical location in the memory (the "physical address") in which the required information is stored. In response, the memory transmits back to the microprocessor the information stored in memory location designated by the physical address.

The time required to supply a microprocessor with required information may be decreased by decreasing the time lapse between the memory's receipt of the physical address and the transmission of the required information. This time lapse is determined by the "speed" of the memory. Unfortunately, the cost of memory increases dramatically with the speed of the memory. Thus, it is rarely practical to use the fastest memory available, especially in systems which require large quantities of memory.

Consequently, it has been discovered that by using a relatively small bank of relatively high-speed memory ("cache memory") as a buffer for a larger bank of slower memory ("system memory"), the average information-request to information-supply speed can be greatly improved. Specifically, in a system having a cache memory (a "cache system"), the microprocessor initially requests information it needs from the cache memory. If the information is stored in the cache memory, the request is said to be a "cache hit" and the information is provided to the microprocessor from the cache memory at the faster rate. If the required information is not stored in the cache memory, the information request is said to be a "cache miss" and the information is retrieved from the system memory at the slower transfer rate. When the information is supplied to the microprocessor from the system memory, a copy of the information is typically stored in the cache memory in anticipation of subsequent requests for the same information. The computational efficiency lost due to a cache miss is referred to herein as the "cache miss penalty."

Many modern computer systems are designed to allow software to address "virtual memory" using "virtual addresses", rather than the actual memory on the system using physical addresses. In a system that supports virtual memory (a "virtual memory sys-

tem"), the software is allowed to assume that the process it implements may access the system's entire address space, that the memory is contiguous, and that the memory begins at a particular address.

5 Due to these assumptions, the software need not be aware of the actual configuration or usage of the system's memory. Virtual memory systems and the benefits thereof are described in detail by David A. Patterson and John L. Hennessy in Computer Architecture A Quantitative Approach (1990), published by Morgan Kaufmann Publishers, Inc. San Mateo, California, published by Morgan Kaufmann Publishers, Inc. San Mateo, California, published by Morgan Kaufmann Publishers, Inc. San Mateo, California, pages 432-454.

10 In virtual memory systems, information must still be stored in the actual system memory. Thus, virtual memory systems must provide a mechanism for translating the virtual addresses generated by the software into physical addresses corresponding to actual memory locations of the system memory. One mechanism commonly used to perform this translation is a translation-lookaside buffer. The concept of translation-lookaside buffers is generally well-known in the art, and is described in detail by David A. Patterson and John L. Hennessy in Computer Architecture A Quantitative Approach (1990), pages 437-438.

15 Similar to non-virtual memory systems, virtual memory systems may employ caching techniques to decrease the time required to retrieve required information. Virtual memory systems may be designed with cache memories which supply required information in response to physical addresses ("physical-address-indexed caches"), cache memories which supply required information in response to virtual addresses ("virtual-address-indexed caches"), or both. Virtual-address-indexed caches are described in detail by David A. Patterson and John L. Hennessy in Computer Architecture A Quantitative Approach (1990), at page 460.

20 When information is required in a virtual memory system which employs both a physical-address-indexed cache and a virtual-address-indexed cache, a search for the desired information is performed in the virtual-address-indexed cache while the virtual address is translated to a physical address. If the desired information is not present in the virtual-address-indexed cache, the processor awaits the result of a search for the desired information in the physical-address-indexed cache. Such a search is begun only after the process of determining a physical address is completed. Information retrieval from the physical-address-indexed cache is slower than retrieval from the virtual-address-indexed cache since a translation delay is imposed prior to the search of the physical-address-indexed cache. Thus, a cache miss in the virtual-address-indexed cache results in a cache miss penalty.

In light of the foregoing, in a virtual memory system having a virtual-address-indexed cache and a physical-address-indexed cache, a method and apparatus for reducing the cache miss penalty caused by a cache miss of the virtual-address-indexed cache is clearly desirable. Further, it is clearly desirable to provide an apparatus and method for avoiding the logical-to-physical address translation time for at least some cache information retrievals from the physical-address-indexed cache.

#### SUMMARY OF THE INVENTION

According to one aspect of the present invention, an apparatus for use in a virtual memory system for reducing the delay between generation of a virtual address by a processor core and transmission to a memory of a physical address corresponding to the virtual address is provided. The generation of the virtual address by the processor core effects a request for required information. The apparatus includes a translation lookaside buffer operatively connected with the processor core and the memory for receiving the virtual address and generating an accurate physical address, and a speculative address generating means operatively connected with the translation lookaside buffer and the memory for generating a speculative address. The speculative address generating means transmits the speculative address to the memory prior to the translation lookaside buffer generating the accurate physical address. The apparatus further includes comparing means, operatively connected with the translation lookaside buffer, the memory and the speculative address generating means, for comparing the accurate physical address with the speculative address. The comparing means transmits the accurate physical address to the memory when the accurate physical address does not match the speculative address.

According to another aspect of the present invention, an apparatus for reducing the time required for a processor core in a virtual memory system to retrieve information from a memory configured to supply information in response to a physical address is provided. The apparatus generally includes accurate physical address generating means, operatively connected with the processor core and the memory, for receiving a virtual address corresponding to an accurate physical address and for generating the accurate physical address. The apparatus further includes speculative address generating means, operatively connected with the memory, for generating a speculative physical address prior to the accurate physical address generating means generating the accurate physical address. The speculative address generating means transmits the speculative physical address to the memory to initiate a first lookup operation in the memory based on the speculative physical address.

The apparatus further includes comparing means operatively connected with the accurate physical address generating means, the memory and the speculative address generating means for receiving the accurate physical address from the accurate physical address generating means and receiving the speculative physical address from the speculative address generating means. The comparing means compares the accurate physical address with the speculative physical address and transmits the accurate physical address to the memory to initiate a second lookup operation based on the accurate physical address when the accurate physical address is different from the speculative physical address.

According to another aspect of the invention, a logical-to-physical address translation unit for transmitting a physical address over a bus in response to a virtual address is provided. The logical-to-physical address translation unit includes accurate translation means for receiving the virtual address and translating the virtual address to an accurate physical address, and speculative translation means, operatively connected with the accurate translation means, for generating a speculative physical address. The logical-to-physical address translation unit further includes comparing means operatively connected with the accurate translation means, the speculative translation means and the bus, for transmitting the speculative physical address as the physical address over the bus prior to completion of the generation of the accurate physical address, for comparing the accurate physical address with the speculative physical address, and for aborting transmission of the speculative physical address over the bus and initiating transmission of the accurate physical address over the bus when the accurate physical address does not match the speculative physical address.

According to yet another aspect of the invention, a virtual memory system is provided. The virtual memory system includes a processor core disposed to generate a virtual address representative of required information and a first cache disposed to supply information in response to virtual addresses. The first cache is operatively connected with the processor core for receiving the virtual address from the processor core. The first cache includes a first plurality of storage locations for storing information. The first cache, in response to receipt of the virtual address, transmits the required information to the processor core when the required information resides in the first plurality of storage locations. The virtual memory system further includes a logical-to-physical address translation means, operatively connected with the processor core, for receiving the virtual address, for generating a physical address in response to the virtual address, and for transmitting the physical address, and a second cache disposed to supply information in response to physical addresses. The

second cache is operatively connected with the logical-to-physical translation means and to the processor core, for receiving the physical address from logical-to-physical translation means. The second cache includes a second plurality of storage locations for storing information. The second cache initiates a lookup operation for the required information in response to receipt of the physical address. The second cache transmits the required information to the processor core when the required information resides in the second plurality of storage locations.

The logical-to-physical translation means includes accurate translation means for receiving the virtual address and translating the virtual address to an accurate physical address and speculative translation means, operatively connected with the accurate translation means, for generating a speculative physical address. The logical-to-physical address translation means further includes comparing means operatively connected with the accurate address generating means, the speculative address generating means and the second cache, for transmitting the speculative physical address as the physical address to the second cache prior to completion of the generation of the accurate physical address, for comparing the accurate physical address to the speculative physical address, and for aborting transmission of the speculative physical address to the second cache and initiating transmission of the accurate physical address to the second cache when the accurate physical address does not match the speculative physical address.

According to yet another aspect of the invention, a method for converting a virtual address to a physical address in a system including a processor core configured to generate the virtual address to request required information from a memory is provided. The method includes the steps of initiating a conversion operation to convert the virtual address to an accurate physical address, generating a speculative physical address, transmitting the speculative physical address to the memory prior to completion of the conversion operation, and comparing the accurate physical address to the speculative physical address. The method further comprises the steps of terminating the transmission of the speculative physical address to the memory, and initiating transmission of the accurate physical address to the memory, when the accurate physical address does not match the speculative physical address.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The features of the present invention, are set forth with particularity in the appended claims. The invention, together with further objects and advantages thereof, may best be understood by making reference to the following description taken in conjunction with

the accompanying drawings, in the several figures of which like reference numerals identify like elements, and wherein:

Figure 1 illustrates in block diagram form a virtual memory system according to the preferred embodiment of the invention; and

Figure 2 illustrates in block diagram form the presently preferred embodiment of a logical-to-physical translation unit.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Figure 1 illustrates in block diagram form a virtual memory system 10 according to the preferred embodiment of the invention. Virtual memory system 10 generally includes a processor core 12, a virtual-address-indexed cache memory 14, a physical-address-indexed cache memory 16 and a logical-to-physical translation unit 18.

Processor core 12 generally represents those components in a processor 13 which decode and execute instructions. Processor core 12 is coupled to physical-address-indexed cache memory 16 through a line buffer 20 by a plurality of data buses 22 and 24. Processor core 12 is coupled to logical-to-physical translation unit 18 by a virtual address bus 26 and by a required-information status line 30. Processor core 12 is also coupled to virtual-address-indexed cache memory 14 by virtual address bus 26. Since system 10 is a virtual memory system, the instructions executed by processor core 12 refer to logical addresses rather physical addresses. When information is required by processor core 12, processor core 12 generates a virtual address identifying the logical location of the required information. Processor core 12 further generates a status signal over line 30 indicative of whether the required information is data or an instruction.

Virtual-address-indexed cache memory 14 is internal, residing on processor 13 with processor core 12. Virtual-address-indexed cache memory 14 is coupled to data bus 22 and includes a plurality of storage locations 15. Virtual-address-indexed cache memory 14 is arranged to transmit information stored within its storage locations 15 in response to virtual addresses received from processor core 12 over virtual address bus 26.

Physical-address-indexed cache memory 16 is external, residing exterior to processor 13. Physical-address-indexed cache memory 16 is coupled to logical-to-physical translation unit 18 and line buffer 20 by a physical address bus 32 and a speculative read control line 33. Physical-address-indexed cache memory 16 includes a plurality of storage locations 17 for storing information, and control circuitry for supplying over data bus 24 information stored in storage locations 17. In contrast to virtual-address-in-

dexed cache memory 14, the information within storage locations 17 is supplied by physical-address-indexed cache memory 16 in response to physical addresses.

Since physical-address-indexed cache memory 16 supplies information in response to physical addresses, the virtual addresses generated by processor core 12 must be translated by logical-to-physical translation unit 18 into physical addresses before required information can be retrieved from physical-address-indexed cache memory 16.

The process by which processor core 12 retrieves required information shall now be explained in greater detail. As mentioned above, when processor core 12 requires information, processor core 12 generates a virtual address indicative of the required information. Processor core 12 transmits this virtual address on virtual address bus 26. Both logical-to-physical translation unit 18 and virtual-address-indexed cache memory 14 receive the virtual addresses transmitted by processor core 12 over virtual address bus 26. In response to the receipt of the virtual address, logical-to-physical translation unit 18 begins to translate the virtual address to a physical address. Simultaneously, virtual-address-indexed cache memory 14 determines whether the required information currently resides in any of its storage locations 15. If the required data does reside in one of storage locations 15, virtual-address-indexed cache memory 14 transmits the required information to processor core 12 over data bus 22.

If the required data does not reside in any of storage locations 15, logical-to-physical translation unit 18 completes the translation of the virtual address into a physical address, and generates the physical address to physical-address-indexed cache memory 16 over physical address bus 32. Upon receipt of the physical address, physical-address-indexed cache memory 16 performs a search to determine if the required information resides in storage locations 17. If the required information is found, physical address-based cache memory 16 transmits the information to processor core 12 through line buffer 20 over data buses 24 and 22. Virtual-address-indexed cache memory 14 may optionally be configured to read and store the information sent from physical-address-indexed cache memory 16 to processor core 12 into one of its storage locations 15 in anticipation of later requests for the same information.

Logical-to-physical translation unit 18 generally includes an accurate address generating unit 40 and a speculative address generating unit 42. When logical-to-physical translation unit 18 receives a virtual address on virtual address bus 26, accurate address generating unit 40 and speculative address generating unit 42, in parallel, generate physical addresses.

Accurate address generating unit 40 is configured to generate physical addresses that always ac-

curately correspond to the storage locations indicated by the received virtual addresses ("accurate physical addresses"). Thus, accurate address generating unit 40 may be, for example, a translation lookaside buffer implemented as commonly known in the art.

5 The logical-to-physical address translation process implemented by speculative address generating unit 42 differs from that implemented by accurate address generating unit 40 in two significant ways. First, the process implemented by speculative address generating unit 42 need not guarantee that the generated physical address will be accurate. Second, the process implemented by speculative address generating unit 42 must complete the generation of a speculative address in less time than it takes for accurate address generating unit 40 to generate the accurate physical address. As long as these two criteria are met, the specific embodiment of speculative address generating unit 42 may vary. The presently preferred embodiment of speculative address generating unit 42 is responsive to the status signal on line 30, and will be described in greater detail below.

10 Thus, after logical-to-physical translation unit 18 receives a virtual address, and while accurate address generating unit 40 is determining the accurate physical address, speculative address generating unit 42 completes generation of a speculative address. Logical-to-physical translation unit 18 transmits the speculative address to physical-address-indexed cache memory 16 to initiate a lookup operation based on the speculative address. Logical-to-physical translation unit 18 further transmits a signal over line 33 to indicate that the physical address currently being supplied to physical-address-indexed cache memory 16 is a speculative address. For example, logical-to-physical translation unit 18 may drive line 33 HIGH when the physical address currently being supplied to physical-address-indexed cache memory 16 is a speculative address.

15 20 25 30 35 40 45 50 Accurate address generating unit 40 completes generation of the accurate physical address after the generation of the speculative address, but prior to the completion of the lookup operation based on the speculative address. When accurate address generating unit 40 completes generation of the accurate physical address, logical-to-physical translation unit 18 compares the accurate physical address with the speculative address. If the speculative address is the same as the accurate physical address, then the lookup operation based on the speculative address is allowed to complete. If the required information is found in physical-address-indexed cache memory 16, it is supplied to processor core 12 through line buffer 20 via data buses 24 and 22.

55 On the other hand, if the accurate physical address is different from the speculative address, logical-to-physical translation unit 18 aborts the transmission of the speculative address over physical ad-

dress bus 32, begins transmitting the accurate physical address to physical-address-indexed cache memory 16 over physical address bus 32, and transmits a signal over line 33 indicating that the address on bus 32 is not a speculative address. For example, logical-to-physical translation unit 18 may drive line 33 LOW when the physical address currently being supplied to physical-address-indexed cache memory 16 is not a speculative address.

Because the control signal on line 33 changes state when the transmission of the speculative address is aborted and the transmission of the accurate physical address is initiated, the control signal effectively indicates to memory 16 the outcome of the comparison of the speculative address and the accurate physical address. While the control signal on line 33 has been described herein as one which changes states to indicate the outcome of the comparison of the speculative address and the accurate physical address, the specific signal indication used to indicate the comparison outcome may vary.

Upon receipt of the new address on bus 32 and the change in the signal over line 33, the control circuitry of physical-address-indexed cache memory 16 aborts the lookup operation based on the speculative address and initiates a lookup operation based on the accurate physical address. If the required information is found in the physical-address-indexed cache memory 16, it is supplied to processor core 12 through line buffer 20 via data buses 24 and 22.

When the speculative address is inaccurate, physical-address-indexed cache memory 16 receives the accurate physical address over physical address bus 32 in substantially the same amount of time it would have taken to generate the accurate address in the absence of speculative address generating unit 42. Hence, no additional delay is imposed by the parallel address generation of speculative address generating unit 42. On the other hand, when the speculative address is accurate, the physical-address-indexed cache memory 16 is able to supply processor core 12 with required information faster than is possible in systems which rely exclusively on always-accurate logical-to-physical address translation techniques.

The actual performance increase obtained by operating speculative address generating unit 42 in parallel with accurate address generating unit 40 depends on how quickly and how accurately the process implemented in speculative address generating unit 42 can predict the correct physical address. The presently preferred embodiment of a speculative address generating unit shall now be described in greater detail with reference to Figure 2.

Figure 2 illustrates in block diagram form the presently preferred embodiment of logical-to-physical translation unit 18. Logical-to-physical translation unit 18 is arranged to work with processor core 12 of

a virtual memory system which employs virtual page techniques to implement virtual memory (a "virtual page system"). In virtual page systems, the storage locations in system memory are divided into a plurality of equal-sized groups, called pages. Each page has a physical page address which uniquely identifies the page, and each storage location within each page has a unique offset from its respective physical page address. Thus, information in any given storage

5 location in the system memory may be accessed by sending to the system memory a physical address containing the physical page address and the offset which, together, uniquely identify the designated storage location.

10 In virtual page systems, software references storage locations in the system memory by virtual addresses which include a virtual page address and an offset. The offset component of a virtual address is identical to the offset component of the physical address of the storage location represented by the virtual address. Therefore, the offset portion of virtual addresses need not be changed during the logical-to-physical address translation process. The virtual page address portion of the virtual address, however,

15 must be translated into the physical page address to which it corresponds before information can be retrieved from the system memory.

20 Because processor core 12 belongs to a virtual page system, virtual addresses generated by processor core 12 include a logical page address and an offset. Virtual address bus 26 includes a logical page address bus 58 over which the logical page address portion of each virtual address is transmitted to memory 16, and an offset bus 60 over which the offset portion of each virtual address is transmitted to memory 16.

25

30

35

40

45

50

55

60

65

70

75

80

85

90

95

100

105

110

115

120

125

130

135

140

145

150

155

160

165

170

175

180

185

190

195

200

205

210

215

220

225

230

235

240

245

250

255

260

265

270

275

280

285

290

295

300

305

310

315

320

325

330

335

340

345

350

355

360

365

370

375

380

385

390

395

400

405

410

415

420

425

430

435

440

445

450

455

460

465

470

475

480

485

490

495

500

505

510

515

520

525

530

535

540

545

550

555

560

565

570

575

580

585

590

595

600

605

610

615

620

625

630

635

640

645

650

655

660

665

670

675

680

685

690

695

700

705

710

715

720

725

730

735

740

745

750

755

760

765

770

775

780

785

790

795

800

805

810

815

820

825

830

835

840

845

850

855

860

865

870

875

880

885

890

895

900

905

910

915

920

925

930

935

940

945

950

955

960

965

970

975

980

985

990

995

1000

1005

1010

1015

1020

1025

1030

1035

1040

1045

1050

1055

1060

1065

1070

1075

1080

1085

1090

1095

1100

1105

1110

1115

1120

1125

1130

1135

1140

1145

1150

1155

1160

1165

1170

1175

1180

1185

1190

1195

1200

1205

1210

1215

1220

1225

1230

1235

1240

1245

1250

1255

1260

1265

1270

1275

1280

1285

1290

1295

1300

1305

1310

1315

1320

1325

1330

1335

1340

1345

1350

1355

1360

1365

1370

1375

1380

1385

1390

1395

1400

1405

1410

1415

1420

1425

1430

1435

1440

1445

1450

1455

1460

1465

1470

1475

1480

1485

1490

1495

1500

1505

1510

1515

1520

1525

1530

1535

1540

1545

1550

1555

1560

1565

1570

1575

1580

1585

1590

1595

1600

1605

1610

1615

1620

1625

1630

1635

1640

1645

1650

1655

1660

1665

1670

1675

1680

1685

1690

1695

1700

1705

1710

1715

1720

1725

1730

1735

1740

1745

1750

1755

1760

1765

1770

1775

1780

1785

1790

1795

1800

1805

1810

1815

1820

1825

1830

1835

1840

1845

1850

1855

1860

1865

1870

1875

1880

1885

1890

1895

1900

1905

1910

1915

1920

1925

1930

1935

1940

1945

1950

1955

1960

1965

1970

1975

1980

1985

1990

1995

2000

2005

2010

2015

2020

2025

2030

2035

2040

2045

2050

2055

2060

2065

2070

2075

2080

2085

2090

2095

2100

2105

2110

2115

2120

2125

2130

2135

2140

2145

2150

2155

2160

2165

2170

2175

2180

2185

2190

2195

2200

2205

2210

2215

2220

2225

2230

2235

2240

2245

2250

2255

2260

2265

2270

2275

2280

2285

2290

2295

2300

2305

2310

2315

2320

2325

2330

2335

2340

2345

2350

2355

2360

2365

2370

2375

2380

2385

2390

2395

2400

2405

2410

2415

2420

2425

2430

2435

2440

2445

2450

2455

2460

2465

2470

2475

2480

2485

2490

2495

2500

2505

2510

2515

2520

2525

2530

2535

2540

2545

2550

2555

2560

2565

2570

2575

2580

2585

2590

2595

2600

2605

2610

2615

2620

2625

2630

2635

2640

2645

2650

2655

2660

2665

2670

2675

2680

2685

2690

2695

2700

2705

2710

2715

2720

2725

2730

2735

2740

2745

2750

2755

2760

2765

2770

2775

2780

2785

2790

2795

2800

2805

2810

2815

2820

2825

2830

2835

2840

2845

2850

2855

2860

2865

2870

2875

2880

2885

2890

2895

2900

2905

2910

2915

2920

2925

2930

2935

2940

2945

2950

2955

2960

2965

2970

2975

2980

2985

2990

2995

3000

3005

3010

3015

3020

3025

3030

3035

3040

3045

3050

3055

3060

3065

3070

3075

3080

3085

3090

3095

3100

3105

3110

3115

3120

3125

3130

3135

3140

3145

3150

3155

3160

3165

3170

3175

3180

3185

3190

3195

3200

3205

3210

3215

3220

3225

3230

3235

3240

3245

3250

3255

3260

3265

3270

3275

3280

3285

3290

3295

3300

3305

3310

3315

3320

3325

3330

3335

3340

3345

3350

3355

3360

3365

3370

3375

3380

3385

3390

3395

3400

3405

3410

3415

3420

3425

3430

3435

3440

3445

3450

3455

3460

3465

3470

3475

3480

3485

3490

3495

3500

3505

3510

3515

3520

3525

3530

3535

3540

3545

3550

3555

3560

3565

3570

3575

3580

3585

3590

3595

3600

3605

3610

3615

3620

3625

3630

3635

3640

3645

3650

3655

3660

3665

3670

3675

3680

3685

3690

3695

3700

3705

3710

3715

3720

3725

3730

3735

3740

3745

3750

3755

3760

3765

3770

3775

3780

3785

3790

3795

3800

3805

3810

3815

3820

3825

3830

3835

3840

3845

3850

3855

3860

3865

3870

3875

3880

3885

3890

3895

3900

3905

3910

3915

3920

3925

3930

3935

3940

3945

3950

3955

3960

3965

3970

3975

3980

3985

3990

3995

4000

4005

4010

4015

4020

4025

4030

4035

4040

4045

4050

4055

4060

4065

4070

4075

4080

4085

4090

4095

4100

4105

4110

4115

4120

4125

4130

4135

4140

4145

4150

4155

4160

4165

4170

4175

4180

4185

4190

4195

4200

4205

4210

4215

4220

4225

4230

4235

4240

4245

4250

4255

4260

4265

4270

4275

4280

4285

4290

4295

4300

4305

4310

4315

4320

4325

4330

4335

4340

4345

4350

4355

4360

4365

4370

4375

4380

4385

4390

4395

4400

4405

4410

4415

4420

4425

4430

4435

4440

4445

4450

4455

4460

4465

4470

4475

4480

4485

4490

4495

4500

4505

4510

4515

4520

4525

4530

4535

4540

4545

4550

4555

4560

4565

4570

4575

4580

4585

4590

4595

4600

4605

4610

4615

4620

4625

4630

4635

4640

4645

4650

4655

4660

4665

4670

4675

4680

4685

4690

4695

4700

4705

4710

4715

4720

4725

4730

4735

4740

4745

4750

4755

4760

4765

4770

4775

4780

4785

4790

4795

4800

4805

4810

4815

4820

4825

4830

4835

4840

4845

4850

4855

4860

4865

4870

4875

4880

4885

4890

4895

4900

4905

4910

4915

4920

4925

4930

4935

4940

4945

4950

4955

4960

4965

4970

4975

4980

4985

4990

4995

5000

5005

5010

5015

5020

5025

5030

5035

5040

5045

5050

5055

5060

5065

5070

5075

5080

5085

5090

5095

5100

5105

5110

5115

5120

5125

5130

5135

5140

5145

5150

5155

5160

5165

5170

5175

5180

5185

5190

5195

5200

5205

5210

5215

5220

5225

5230

5235

5240

5245

5250

5255

5260

5265

5270

5275

5280

5285

5290

5295

5300

5305

5310

5315

5320

5325

5330

5335

5340

5345

5350

5355

5360

5365

5370

5375

5380

5385

5390

5395

5400

5405

5410

5415

5420

5425

5430

5435

5440

5445

5450

5455

5460

5465

5470

5475

5480

5485

5490

5495

5500

5505

5510

5515

5520

5525

5530

5535

5540

5545

5550

5555

5560

5565

5570

5575

5580

5585

5590

5595

5600

5605

5610

5615

5620

5625

5630

5635

5640

5645

5650

5655

5660

5665

5670

5675

5680

5685

5690

5695

5700

5705

5710

5715

5720

ical-to-physical translation unit 18 receives from processor core 12, over virtual address bus 26, a virtual address which identifies the storage location that holds required information. More specifically, the logical page address corresponding to the required information's physical storage location is transmitted on logical page address bus 58, and the offset corresponding to the required information's physical storage location is transmitted on offset bus 60. Processor core 12 also transmits over line 30 a signal indicating whether the required information represents an instruction or data.

In response to the virtual address, translation lookaside buffer 62 begins converting the logical page address of the virtual address to an accurate physical page address. Simultaneously, speculative address generating unit 42 causes a speculative address to be sent to memory 16 over physical page address bus 58 via comparing unit 68 and a bus 74, and control logic 84 transmits a signal to memory 16 on line 33 to indicate that the physical address on bus 32 is a speculative address. In response to receipt of the speculative address on bus 32 and the signal on line 33, memory 16 initiates a lookup operation based on the speculative address.

When translation lookaside buffer 62 finishes translating the logical page address to an accurate physical page address, translation lookaside buffer 62 transmits the accurate physical page address to the comparing unit 68 over a bus 72. Comparing unit 68 compares the accurate physical page address with the speculative page address. If the accurate physical page address matches the speculative page address, then the comparing unit 68 completes the transmission of the speculative page address over physical page address bus 58 and memory 16 completes the lookup operation based on the speculative address.

If the accurate physical page address does not match the speculative page address, then comparing unit 68 ceases to transmit the speculative page address on physical page address bus 58, sends a signal to control logic 84 over line 95 to indicate that the current speculative page address is not accurate, and begins to transmit the accurate physical page address over physical page address bus 58.

In response to the signal on line 95, control logic 84 transmits a signal on line 33 to indicate that the physical address on bus 32 is no longer a speculative address. Upon receipt of the signal on line 33, memory 16 aborts the lookup operation based on the speculative address and initiates a search based on the accurate physical address.

Speculative address generating unit 42 is configured to generate a speculative page address based on the accurate physical page address most recently generated by translation lookaside buffer 62 for information of the same type as the information that is currently required. For example, if, in response to the last

request for an instruction, the translation lookaside buffer 62 generated a physical page address 0, then speculative address generating unit 42 will generate, in response to a request for the retrieval of an instruction, page address 0 as a speculative page address. Similarly, if, in response to the last request for data, the translation lookaside buffer 62 generated physical page address 1, then speculative address generating unit 42 will generate, in response to a request for the retrieval of data, page address 1 as a speculative page address.

Speculative address generating unit 42 implements the above prediction process as follows. After translation lookaside buffer 62 generates an accurate physical page address, translation lookaside buffer 62 transmits the accurate physical page address over bus 73 to instruction address latch 86 and data address latch 88. In response to the signal on line 30, control logic 84 determines whether the accurate physical page address was generated in response to a request for an instruction or a request for data. If the accurate physical page address was generated in response to a request for an instruction, control logic 84 transmits a control signal over line 92 to instruction address latch 86 to cause the accurate physical page address on bus 73 to be latched into instruction address latch 86. On the other hand, if the accurate physical page address was generated in response to a request for data, control logic 84 transmits a control signal over line 94 to data address latch 88 to cause the accurate physical page address on bus 73 to be latched into data address latch 88.

In response to a subsequent request for information, control logic 84 transmits a signal over line 96 to multiplexer 90 to select either instruction address latch 86 or data address latch 88. Specifically, if the signal on line 30 indicates that an instruction is being requested by processor core 12, then control logic 84 transmits a signal over line 96 to multiplexer 90 to cause the page address previously stored in instruction address latch 86 to be transmitted through multiplexer 90 onto bus 74. If the signal on line 30 indicates that data is being requested by processor core 12, control logic 84 transmits a signal over line 96 to multiplexer 90 to cause the page address stored in instruction address latch 86 to be transmitted through multiplexer 90 onto bus 74.

Thus, the preferred embodiment of speculative address generating unit 42 stores the most recently generated actual page address for an instruction request in instruction address latch 86, and the most recently generated actual page address for a data request in data address latch 88. When a request for an instruction is received, speculative address generating unit 42 immediately transmits the contents of instruction address latch 86 as a speculative page address. Likewise, when a request for data is received, speculative address generating unit 42 immediately

transmits the contents of data address latch 88 as a speculative page address.

The physical page of the currently required data or instruction is not necessarily the physical page of the most recently requested data or instruction. For this reason, the speculative page address must always be checked against an accurate page address. However, related instructions tend to be clustered together in memory, and related data tend to be clustered together in memory. Further, to perform any given operation, a processor is generally required to execute a plurality of related instructions on a plurality of related data. Consequently, there tends to be a high correlation between the physical page of data requested and the physical page of data previously requested. Likewise, there tends to be a high correlation between the physical page of instructions requested and the physical page of instructions previously requested. The stronger these correlations, the more often the speculative address generated by speculative address generating unit 42 is accurate.

As can be seen from the foregoing, an apparatus and method are provided which reduce the average time required to retrieve information from a physical-address-indexed cache in a virtual memory system. The access time is reduced by initiating a preliminary search for the required information while the virtual address generated by the processor core is being converted to an accurate physical address. The preliminary search is based on a speculative address, which, in the preferred embodiment of the invention, is the physical address of the most recently retrieved information of the same type. Once an accurate physical address has been determined, it is compared with the speculative address. If the speculative and the actual addresses match, then the preliminary search is allowed to complete. If the speculative and the actual addresses do not match, the preliminary search is aborted and a search is begun based on the accurate physical address.

It is to be understood that the detailed drawings and specific examples given describe a preferred embodiment of the invention and are for the purpose of illustration only, and that the apparatus of the invention is not limited to the precise details and conditions disclosed. For example, while an internal virtual-address-indexed cache memory and an external physical-address-indexed cache memory have been described, these caches may be either internal or external. In addition, while the present invention has been described with reference to a physical-address-indexed cache, it may alternatively be used with any physical-address-indexed memory, such as system memory. Various other changes may also be made without departing from the spirit of the invention which is defined by the following claims:

The following clauses set out various aspects of the invention.

1. An apparatus for use in a virtual memory system for reducing the delay between generation of a virtual address by a processor core and transmission to a memory of a physical address corresponding to said virtual address, said generation of said virtual address by said processor core effecting a request for required information, the apparatus comprising:

a translation lookaside buffer operatively connected with said processor core and said memory for receiving said virtual address and generating an accurate physical address;

a speculative address generating means operatively connected with said translation lookaside buffer and said memory for generating a speculative address, said speculative address generating means transmitting said speculative address to said memory prior to said translation lookaside buffer generating said accurate physical address;

comparing means, operatively connected with said translation lookaside buffer, said memory and said speculative address generating means, for comparing said accurate physical address with said speculative address;

said comparing means transmitting said accurate physical address to said memory when said accurate physical address does not match said speculative address.

2. An apparatus for use in a virtual memory system as recited in clause 1 wherein said memory is configured to initiate a first lookup operation based on said speculative address substantially upon receipt of said speculative address, said speculative address generating means being configured to transmit to said memory a control signal indicative of whether said speculative address matches said accurate physical address, said memory being configured to abort said first lookup operation based on said speculative address and initiate a second lookup operation based on said accurate physical address when said control signal indicates that said speculative address does not match said accurate physical address.

3. An apparatus for use in a virtual memory system as recited in clause 1 wherein said speculative address generating means includes storage means for storing a previous physical address, said previous physical address being a selected accurate physical address previously generated by said translation lookaside buffer, at least a portion of said speculative address being related with at least a portion of said previous physical address according to a predetermined relationship.

4. An apparatus for use in a virtual memory system as recited in clause 3 wherein said virtual

memory system includes a virtual-address-indexed cache memory operatively connected with said processor core, said virtual-address-indexed cache memory receiving said virtual address from said processor core and transmitting said required information to said processor core when said required information is present in said virtual-address-indexed cache memory.

5. An apparatus for use in a virtual memory system as recited in clause 1 wherein said generation of said virtual address by said processor core effects a present request for said required information, said present request being a present instruction request when said required information is an instruction, said present request being a present data request when said required information is data, said speculative address generating means including storage means for storing at least one previous instruction address and at least one previous data address, said at least one previous instruction address being a first selected accurate physical address previously generated by said translation lookaside buffer, said at least one previous data address being a second selected accurate physical address previously generated by said translation lookaside buffer, at least a portion of said speculative address being related with at least a portion of said at least one previous instruction address according to a predetermined relationship when said present request is said present instruction request, at least a portion of said speculative address being related with at least a portion of said at least one previous data address according to a predetermined relationship when said present request is said present data request.

6. An apparatus for use in a virtual memory system as recited in clause 1 wherein said memory is a cache memory configured to supply said required information in response to said physical address.

7. An apparatus for reducing the time required for a processor core in a virtual memory system to retrieve information from a memory configured to supply information in response to a physical address, the apparatus comprising:

accurate physical address generating means, operatively connected with said processor core and said memory, for receiving a virtual address corresponding to an accurate physical address and for generating said accurate physical address;

speculative address generating means, operatively connected with said memory, for generating a speculative physical address prior to said accurate physical address generating means generating said accurate physical address, said speculative address generating

means transmitting said speculative physical address to said memory to initiate a first lookup operation in said memory based on said speculative physical address;

comparing means operatively connected with said accurate physical address generating means, said memory and said speculative address generating means for receiving said accurate physical address from said accurate physical address generating means and receiving said speculative physical address from said speculative address generating means, said comparing means comparing said accurate physical address with said speculative physical address;

said comparing means transmitting said accurate physical address to said memory to initiate a second lookup operation based on said accurate physical address when said accurate physical address is different from said speculative physical address.

8. An apparatus for reducing the time required for a processor core in a virtual memory system to retrieve information from a memory configured to supply information in response to a physical address, as recited in clause 7, wherein said memory is configured to initiate a first lookup operation based on said speculative physical address substantially upon receipt of said speculative physical address, said speculative address generating means being configured to transmit a control signal to said memory, said control signal being indicative of whether said speculative physical address matches said accurate physical address, said memory being configured to abort said first lookup operation based on said speculative physical address and initiate a second lookup operation based on said accurate physical address when said control signal indicates that said speculative physical address does not match said accurate physical address.

9. An apparatus for reducing the time required for a processor core in a virtual memory system to retrieve information from a memory configured to supply information in response to a physical address, as recited in clause 7, wherein said speculative address generating means is operatively connected with said accurate physical address generating means, said speculative address generating means including storage means for storing at least one previous physical address, said at least one previous physical address being a selected accurate physical address previously generated by said accurate physical address generating means, at least a portion of said speculative physical address being related with at least a portion of said at least one previous physical address according to a predetermined relationship.

10. An apparatus for reducing the time required

for a processor core in a virtual memory system to retrieve information from a memory configured to supply information in response to a physical address, as recited in clause 9, wherein said virtual memory system includes a virtual-address-indexed cache memory operatively connected with said processor core, said processor core generating said virtual address to effect a request for required information, said virtual-address-indexed cache memory receiving said virtual address from said processor core, said virtual-address-indexed cache memory transmitting said required information to said processor core when said required information is present in said virtual-address-indexed cache memory.

11. An apparatus for reducing the time required for a processor core in a virtual memory system to retrieve information from a memory configured to supply information in response to a physical address, as recited in clause 7, wherein said processor core generates said virtual address to effect a present request for said information, said present request being a present instruction request when said information is an instruction, said present request being a present data request when said information is data, said speculative address generating means including storage means for storing at least one previous instruction address and at least one previous data address, said at least one previous instruction address being an accurate physical address previously generated by said accurate physical address generating means, said at least one previous data address being an accurate physical address previously generated by said accurate physical address generating means in response to said at least one previous data request, at least a portion of said speculative physical address being related with at least a portion of said at least one previous instruction address according to a predetermined relationship when said present request is said present instruction request, at least a portion of said speculative physical address being related with at least a portion of said at least one previous data address according to a predetermined relationship when said present request is said present data request.

12. A logical-to-physical address translation unit for transmitting a physical address over a bus in response to a virtual address, comprising:

accurate translation means for receiving said virtual address and translating said virtual address to an accurate physical address;

**speculative translation means, operatively connected with said accurate translation means, for generating a speculative physical address;**

comparing means operatively connected

with said accurate translation means, said speculative translation means and said bus, for transmitting said speculative physical address as said physical address over said bus prior to completion of said generation of said accurate physical address, for comparing said accurate physical address with said speculative physical address, and for aborting transmission of said speculative physical address over said bus and initiating transmission of said accurate physical address over said bus when said accurate physical address does not match said speculative physical address.

13. A logical-to-physical address translation unit for transmitting a physical address in response to a virtual address, as recited in clause 12, further comprising a memory operatively connected with said bus, said memory being configured to initiate a first lookup operation based on said speculative physical address substantially upon receipt of said speculative physical address over said bus, said speculative translation means being configured to transmit a control signal to said memory, said control signal being indicative of whether said speculative physical address matches said accurate physical address, said memory being configured to abort said first lookup operation based on said speculative physical address and initiate a second lookup operation based on said accurate physical address when said control signal indicates that said speculative physical address does not match said accurate physical address.

14. A logical-to-physical address translation unit for transmitting a physical address in response to a virtual address, as recited in clause 12, wherein said accurate translation means is a translation lookaside buffer.

15. A logical-to-physical address translation unit for transmitting a physical address in response to a virtual address, as recited in clause 12, wherein said virtual address comprises a virtual page portion and an offset portion, said accurate translation means generating an accurate physical page address, said accurate physical address comprising said accurate physical page address and said offset portion, said speculative translation means generating a speculative page address, said speculative physical address comprising said speculative page address and said offset portion.

16. A logical-to-physical address translation unit for transmitting a physical address in response to a virtual address, as recited in clause 12, wherein said speculative translation means includes storage means for storing at least one previous physical address, said at least one previous physical address being a selected accurate physical ad-

dress previously generated by said accurate translation means, at least a portion of said speculative physical address being related with said at least one previous physical address according to a predetermined relationship.

17. A logical-to-physical address translation unit for transmitting a physical address in response to a virtual address, as recited in clause 12, wherein said virtual address represents a present request for required information, said present request being a present instruction request when said required information is an instruction, said present request being a present data request when said required information is data, said speculative translation means including storage means for storing at least one previous instruction address and at least one previous data address, said at least one previous instruction address being a first selected accurate physical address previously generated by said accurate translation means, said at least one previous data address being a second selected accurate physical address previously generated by said accurate translation means, at least a portion of said speculative physical address being related with at least a portion of said at least one previous instruction address according to a predetermined relationship when said present request is said present instruction request, at least a portion of said speculative physical address being related with at least a portion of said at least one previous data address according to a predetermined relationship when said present request is said present data request.

18. A virtual memory system, including:

a processor core disposed to generate a virtual address representative of required information;

a first cache disposed to supply information in response to virtual addresses, said first cache being operatively connected with said processor core for receiving said virtual address from said processor core, said first cache including a first plurality of storage locations for storing information, said first cache transmitting said required information to said processor core in response to receipt of said virtual address when said required information resides in said first plurality of storage locations;

a logical-to-physical address translation means, operatively connected with said processor core, for receiving said virtual address, for generating a physical address in response to said virtual address, and for transmitting said physical address; and

a second cache disposed to supply information in response to physical addresses, said second cache being operatively connected with

5 said logical-to-physical address translation means and to said processor core, for receiving said physical address from logical-to-physical address translation means, said second cache including a second plurality of storage locations for storing information, said second cache initiating a lookup operation for said required information in response to receipt of said physical address, said second cache transmitting said required information to said processor core when said required information resides in said second plurality of storage locations;

10 said logical-to-physical address translation means including:

15 accurate translation means for receiving said virtual address and for translating said virtual address to an accurate physical address,

20 speculative translation means, operatively connected with said accurate translation means, for generating a speculative physical address,

25 comparing means operatively connected with said accurate address generating means, said speculative address generating means and said second cache, for transmitting said speculative physical address as said physical address to said second cache prior to completion of said generation of said accurate physical address, for comparing said accurate physical address to said speculative physical address, and for aborting transmission of said speculative physical address to said second cache and initiating transmission of said accurate physical address to said second cache when said accurate physical address does not match said speculative physical address.

30 19. A virtual memory system, as recited in clause 18, wherein said second cache is configured to initiate a first lookup operation based on said speculative physical address substantially upon receipt of said speculative physical address, said speculative translation means being configured to transmit a control signal to said second cache, said control signal being indicative of whether said speculative physical address matches said accurate physical address, said second cache being configured to abort said first lookup operation based on said speculative physical address and initiate a second lookup operation based on said accurate physical address when said control signal indicates that said speculative physical address does not match said accurate physical address.

35 20. A virtual memory system, as recited in of clause 18, wherein said accurate translation means is a translation lookaside buffer.

40 21. A virtual memory system, as recited in clause

18, wherein said virtual address comprises a virtual page portion and an offset portion, said accurate translation means generating an accurate physical page address, said accurate physical address comprising said accurate physical page address and said offset portion, said speculative translation means generating a speculative page address, said speculative physical address comprising said speculative page address and said offset portion.

22. A virtual memory system, as recited in clause 18, wherein said speculative translation means includes storage means for storing at least one previous physical address, said at least one previous physical address being a selected accurate physical address previously generated by said accurate translation means, at least a portion of said speculative physical address being related with at least a portion of said at least one previous physical address according to a predetermined relationship.

23. A virtual memory system, as recited in clause 18, wherein said generation of said virtual address by said processor core effects a present request for said required information, said present request being a present instruction request when said required information is an instruction, said present request being a present data request when said required information is data, said speculative translation means including storage means for storing at least one previous instruction address and at least one previous data address, said at least one previous instruction address being a first selected accurate physical address previously generated by said accurate translation means, said at least one previous data address being a second selected accurate physical address previously generated by said accurate translation means, at least a portion of said speculative physical address being related with at least a portion of said at least one previous instruction address according to a predetermined relationship when said present request is said present instruction request, at least a portion of said speculative physical address being related with at least a portion of said at least one previous data address according to a predetermined relationship when said present request is said present data request.

24. A virtual memory system, as recited in clause 18, wherein said processor core, said logical-to-physical address translation means, and said first cache reside on a single processor chip.

25. A virtual memory system, as recited in clause 24, wherein said second cache resides external to said single processor chip.

26. A method for converting a virtual address to a physical address in a system including a proc-

essor core configured to generate said virtual address to request required information from a memory, comprising the steps of:

5 initiating a conversion operation to convert said virtual address to an accurate physical address;

generating a speculative physical address;

transmitting said speculative physical address to said memory prior to completion of said conversion operation;

comparing said accurate physical address to said speculative physical address; and

10 terminating said transmission of said speculative physical address to said memory, and initiating transmission of said accurate physical address to said memory, when said accurate physical address does not match said speculative physical address.

15 27. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 26, including the further steps of causing said memory to initiate a first lookup operation based on said speculative physical address substantially upon receipt of said speculative physical address, and causing said memory to abort said first lookup operation based on said speculative physical address and initiate a second lookup operation based on said accurate physical address when said speculative physical address does not match said accurate physical address.

20 28. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 26, wherein said conversion operation is performed by a lookaside translation buffer.

25 29. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 26, wherein said virtual address has a virtual page portion and an offset portion, said conversion operation including converting said virtual page portion to an accurate physical page address.

30 30. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 29, wherein said step of generating includes generating a speculative physical page address.

35 31. A method for converting a virtual address to

a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 30, wherein said step of transmitting said speculative physical address includes transmitting said speculative physical page address and said offset portion.

32. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 31, wherein said step of transmitting said accurate physical address includes transmitting said accurate physical page address and said offset portion.

33. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 26, wherein said processor core further generates a status signal representative of whether said required information represents an instruction or data, the method further including the step of receiving said status signal, said generation of said speculative page address being responsive to said status signal.

34. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 33, wherein said generation of said speculative physical address is further responsive to a previously-generated accurate physical address.

35. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in clause 34, said virtual address effecting a present request for said required information, said present request being a present data request when said required information is data, said present request being a present instruction request when said required information is an instruction, wherein said step of generating said speculative physical address includes the steps of:

storing a first previously-generated accurate physical address, said first previously-generated accurate physical address being a first selected accurate physical address generated previous to said accurate physical address in response to a previous data request;

storing a second previously-generated accurate physical address, said second previously-generated accurate physical address being a second selected accurate physical address generated previous to said accurate physical address

in response to a previous instruction request; generating said first previously-generated accurate physical address when said status signal indicates said required information represents data; and

generating said second previously-generated accurate physical address when said status signal indicates said required information represents an instruction.

## Claims

1. An apparatus for use in a virtual memory system for reducing the delay between generation of a virtual address by a processor core and transmission to a memory of a physical address corresponding to said virtual address, said generation of said virtual address by said processor core effecting a request for required information, the apparatus comprising:

a translation lookaside buffer operatively connected with said processor core and said memory for receiving said virtual address and generating an accurate physical address;

a speculative address generating means operatively connected with said translation lookaside buffer and said memory for generating a speculative address, said speculative address generating means transmitting said speculative address to said memory prior to said translation lookaside buffer generating said accurate physical address;

comparing means, operatively connected with said translation lookaside buffer, said memory and said speculative address generating means, for comparing said accurate physical address with said speculative address;

said comparing means transmitting said accurate physical address to said memory when said accurate physical address does not match said speculative address.

2. An apparatus for use in a virtual memory system as recited in Claim 1 wherein said memory is configured to initiate a first lookup operation based on said speculative address substantially upon receipt of said speculative address, said speculative address generating means being configured to transmit to said memory a control signal indicative of whether said speculative address matches said accurate physical address, said memory being configured to abort said first lookup operation based on said speculative address and initiate a second lookup operation based on said accurate physical address when said control signal indicates that said speculative address does not match said accurate physical address.

3. An apparatus for use in a virtual memory system as recited in Claim 1 wherein said speculative address generating means includes storage means for storing a previous physical address, said previous physical address being a selected accurate physical address previously generated by said translation lookaside buffer, at least a portion of said speculative address being related with at least a portion of said previous physical address according to a predetermined relationship.

4. An apparatus for use in a virtual memory system as recited in Claim 3 wherein said virtual memory system includes a virtual-address-indexed cache memory operatively connected with said processor core, said virtual-address-indexed cache memory receiving said virtual address from said processor core and transmitting said required information to said processor core when said required information is present in said virtual-address-indexed cache memory.

5. An apparatus for use in a virtual memory system as recited in Claim 1 wherein said generation of said virtual address by said processor core effects a present request for said required information, said present request being a present instruction request when said required information is an instruction, said present request being a present data request when said required information is data, said speculative address generating means including storage means for storing at least one previous instruction address and at least one previous data address, said at least one previous instruction address being a first selected accurate physical address previously generated by said translation lookaside buffer, said at least one previous data address being a second selected accurate physical address previously generated by said translation lookaside buffer, at least a portion of said speculative address being related with at least a portion of said at least one previous instruction address according to a predetermined relationship when said present request is said present instruction request, at least a portion of said speculative address being related with at least a portion of said at least one previous data address according to a predetermined relationship when said present request is said present data request.

6. An apparatus for use in a virtual memory system as recited in Claim 1 wherein said memory is a cache memory configured to supply said required information in response to said physical address.

7. An apparatus for reducing the time required for a processor core in a virtual memory system to re-

5 retrieve information from a memory configured to supply information in response to a physical address, the apparatus comprising:

10 accurate physical address generating means, operatively connected with said processor core and said memory, for receiving a virtual address corresponding to an accurate physical address and for generating said accurate physical address;

15 speculative address generating means, operatively connected with said memory, for generating a speculative physical address prior to said accurate physical address generating means generating said accurate physical address, said speculative address generating means transmitting said speculative physical address to said memory to initiate a first lookup operation in said memory based on said speculative physical address;

20 comparing means operatively connected with said accurate physical address generating means, said memory and said speculative address generating means for receiving said accurate physical address from said accurate physical address generating means and receiving said speculative physical address from said speculative address generating means, said comparing means comparing said accurate physical address with said speculative physical address;

25 said comparing means transmitting said accurate physical address to said memory to initiate a second lookup operation based on said accurate physical address when said accurate physical address is different from said speculative physical address.

30

35

40

45

50

55

8. A logical-to-physical address translation unit for transmitting a physical address over a bus in response to a virtual address, comprising:

accurate translation means for receiving said virtual address and translating said virtual address to an accurate physical address;

speculative translation means, operatively connected with said accurate translation means, for generating a speculative physical address;

comparing means operatively connected with said accurate translation means, said speculative translation means and said bus, for transmitting said speculative physical address as said physical address over said bus prior to completion of said generation of said accurate physical address, for comparing said accurate physical address with said speculative physical address, and for aborting transmission of said speculative physical address over said bus and initiating transmission of said accurate physical address over said bus when said accurate physical ad-

dress does not match said speculative physical address.

9. A virtual memory system, including:

a processor core disposed to generate a virtual address representative of required information;

a first cache disposed to supply information in response to virtual addresses, said first cache being operatively connected with said processor core for receiving said virtual address from said processor core, said first cache including a first plurality of storage locations for storing information, said first cache transmitting said required information to said processor core in response to receipt of said virtual address when said required information resides in said first plurality of storage locations;

a logical-to-physical address translation means, operatively connected with said processor core, for receiving said virtual address, for generating a physical address in response to said virtual address, and for transmitting said physical address; and

a second cache disposed to supply information in response to physical addresses, said second cache being operatively connected with said logical-to-physical address translation means and to said processor core, for receiving said physical address from logical-to-physical address translation means, said second cache including a second plurality of storage locations for storing information, said second cache initiating a lookup operation for said required information in response to receipt of said physical address, said second cache transmitting said required information to said processor core when said required information resides in said second plurality of storage locations;

said logical-to-physical address translation means including:

accurate translation means for receiving said virtual address and for translating said virtual address to an accurate physical address,

speculative translation means, operatively connected with said accurate translation means, for generating a speculative physical address,

comparing means operatively connected with said accurate address generating means, said speculative address generating means and said second cache, for transmitting said speculative physical address as said physical address to said second cache prior to completion of said generation of said accurate physical address, for comparing said accurate physical address to said speculative physical address,

and for aborting transmission of said speculative physical address to said second cache and initiating transmission of said accurate physical address to said second cache when said accurate physical address does not match said speculative physical address.

10. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, comprising the steps of:

initiating a conversion operation to convert said virtual address to an accurate physical address;

generating a speculative physical address;

transmitting said speculative physical address to said memory prior to completion of said conversion operation;

comparing said accurate physical address to said speculative physical address; and

terminating said transmission of said speculative physical address to said memory, and initiating transmission of said accurate physical address to said memory, when said accurate physical address does not match said speculative physical address.

11. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in Claim 10, including the further steps of causing said memory to initiate a first lookup operation based on said speculative physical address substantially upon receipt of said speculative physical address, and causing said memory to abort said first lookup operation based on said speculative physical address and initiate a second lookup operation based on said accurate physical address when said speculative physical address does not match said accurate physical address.

12. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in Claim 10, wherein said conversion operation is performed by a lookaside translation buffer.

13. A method for converting a virtual address to a physical address in a system including a processor core configured to generate said virtual address to request required information from a memory, as recited in Claim 10, wherein said vir-

tual address has a virtual page portion and an offset portion, said conversion operation including converting said virtual page portion to an accurate physical page address.

5

10

15

20

25

30

35

40

45

50

55

16




FIG. 1




FIG. 2